Astronomy & Astrophysics manuscript no. ms
October 18, 2012
© ESO 2012

Spatially coupled inversion of spectro-polarimetric image data I: Method and first results

M. van Noort

Max-Planck Institute for Solar System Research, Max-Planck Straße 1, D-37412 Katlenburg-Lindau, Germany

October 18, 2012

arXiv:1210.4636v1 [astro-ph.IM] 17 Oct 2012

ABSTRACT

Context. When inverting solar spectra, image degradation effects that are present in the data are usually approximated or not considered.
Aims. We develop a data reduction method that takes these issues into account and minimizes the resulting errors.
Methods. By accounting for the diffraction PSF of the telescope during the inversions, we can produce a self-consistent solution that best fits the observed data, while simultaneously requiring fewer free parameters than conventional approaches.
Results. Simulations using realistic MHD data indicate that the method is stable for all resolutions, including those with pixel scales well beyond those that can be resolved with a 0.5m telescope, such as the Hinode SOT. Application of the presented method to reduce full Stokes data from the Hinode spectro-polarimeter results in dramatically increased image contrast and an increase in the resolution of the data to the diffraction limit of the telescope in almost all Stokes and fit parameters. The resulting data allow for detecting and interpreting solar features that have so far only been observed with 1m class ground-based telescopes.
Conclusions. A new inversion method was developed that allows for accurate fitting of solar spectro-polarimetric imaging data over a large field of view, while simultaneously improving the noise statistics and spatial resolution of the results significantly.

Key words. Techniques: imaging spectroscopy, polarimetric, methods: data analysis, numerical

1. Introduction

In solarphysical research, as in the vast majority of astrophysical research, information about the object under investigation is received in the form of electromagnetic radiation. Since this radiation was emitted by, passed through or reflected off the object of interest, interpreting the information contained in the radiation typically involves computing the observed data from a model of the object using basic physical principles.

Interpreting the data typically starts by making an educated guess about the observed object's atmosphere (the model), followed by an evaluation of the compatibility of the observable, calculated using this model, with the observed data. Additional improvements may then be made, depending on the flexibility of the model. Optimizing the compatibility of the synthesized results with the data is usually conveniently formulated in terms of the minimization of the square of the differences. A number of codes exist that automate this optimization process, either by a standard linearized optimization (SIR (Ruiz Cobo & del Toro Iniesta 1992), SPINOR (Frutiger 2000; Frutiger et al. 2000), NICOLE (Socas-Navarro et al. 1998), HAZEL (Asensio Ramos et al. 2008), and many more); by statistical methods (HELIX (Lagg et al. 2004, 2009), MERLIN, etc.); by means of Bayesian methods (Bayes-ME Asensio Ramos et al. 2007); or by using artificial neural networks (Socas-Navarro 2003; Carroll et al. 2008), with greatly varying levels of detail in their treatment of the physics and often with the aim of treating very specific types of data.

Usually the simplest atmospheric model that can reproduce the observations is used for analysis of the solar lower atmosphere. The model is characterized by a number of nodes, at which the physical parameters that are relevant for the formation of the spectrum are defined. These quantities are then interpolated to a finer grid to allow for the calculation of an accurate solution to the radiative transfer equation. Automatic inversion codes, which optimize this kind of model, have been in use for two decades for interpreting solar data.

Spurred by the increase in computing power in recent times, much effort has gone into increasing the level of detail and self-consistency in the treatment of the basic physics, non-LTE (NICOLE) and scattering polarization effects (HELIX, HAZEL, etc.). Without these, many observed spectral features can be shown to be poorly reproduced, even in cases where the fitted atmosphere is realistic.

The observed data used in this process, however, quantify the properties of radiation that has passed through a telescope, an instrument, and a detector, all of which alter the data depending on the size, optical quality and other instrumental properties. For ground-based observations, in addition, the radiation passes through the earths atmosphere, which redistributes it by refracting and scattering.

A large body of work exists on removing the degradation of the spatial resolution due to an increase in the extent of the point spread function from instrumental and atmospheric effects. In particular the optimization of image information from observed data degraded by atmospheric turbulence ((MO)MFBD (van Noort et al. 2005), phase diversity methods (Paxman et al. 1992; Löfdahl & Scharmer 1994, and references therein), and speckle reconstruction methods (Labeyrie 1970; Knox & Thompson 1974; Lohmann et al. 1983; de Boer et al. 1992; von der Luehe 1993, and many more)) has recently found widespread application in spectral imaging data produced by an array of narrow band filter instruments;

Send offprint requests to: M. v. N.: e-mail: vannoort@mps.mpg.de

e.g. SOUP (Title & Rosenberg 1981), CRISP (Scharmer 2006), GFPI (Puschmann et al. 2006) and IBIS (Cavallini 2006), to name just a few. These methods only work well for variable degradation on timescales [much] shorter than the time resolution of the observation.

In the spectral dimension, on the other hand, degradation by instrumental and atmospheric effects has a rather different effect on the data. Where the spatial degradation merely decreases the resolution in the spatial dimensions, the blending of the spectra results in observed spectra containing a mix of information from a collection of different atmospheric components. Much work has been done on the analysis and interpretation of such "multi-component" spectra (Stenflo 1968; Skumanich & Lites 1991; Bernasconi & Solanki 1996, and many more), by modeling them as arising from two or more separate atmospheric "components", which remain unresolved in the data.

Orozco Suárez et al. (2007) model the effect of telescope diffraction on the spectra as a local "stray-light" contribution, where the light from the surroundings is taking the place of the contribution of a second atmospheric component. The apparent multi-component spectra are not interpreted in terms of unresolved fine structure, but rather in terms of an additive contribution produced by telescope diffraction.

In this paper we develop this idea further and develop a method that rigorously deals with the effect of image degradation by instrumental and telescope effects, by integrating these effects in a 2-D adaptation of the SPINOR inversion code.

2. Inversion of spatially degraded data

Any process that infers the most probable atmospheric structure that gave rise to the observed spectral data may generally be referred to as an inversion process. One could even argue that the inversion is only defined by the way the probability of a particular atmospheric structure is defined, since the method used to locate the point of maximum likelihood is normally quite independent of the formulation of the probability function, but this would not do justice to the information that must be obtained from the atmospheric structure in order to optimize it.

Considerable differences exist between the various inversion codes in the way they describe the atmosphere, the description of the radiative transfer problem, and in the way they try to maximize the likelihood function. However, what they all have in common is that ultimately they all compute synthetic profiles from estimated atmospheres and compare these to the observed profile, in order to infer the most likely atmosphere that produced the observed spectrum.

A physically accurate description is included in most of these codes in ample detail, since it is obvious that this is needed to have any chance of interpreting an observed spectrum in terms of physical quantities. Although the most common assumption made in inversion codes is that of local thermodynamic equilibrium (LTE), there are codes that compute the line profiles in full NLTE (e.g. NICOLE, HAZEL), codes that include detailed scattering physics (HELIX, HAZEL), and codes that include detailed molecular physics (SPINOR). In addition, there is a great variation in the amount of detail in which the atmosphere is parameterized, from slabs with constant properties to full 3-D structures with a full height stratification in all physical quantities (see for instance Asensio Ramos et al. 2012, for an overview).

Instrumental effects, on the other hand, with the exception of spectral smearing, are not generally considered in much detail, if at all, by many codes, because this is considered by most code authors to be taken care of by appropriate data reduction routines. Although careful calibration of the data generally receives a lot of attention, there are always significant limits to the extent to which it is possible to remove the instrumental effects from the data.

In particular, the removal from the data of any degradation of the data that reduces the information content (such as convolution with an extended convolution kernel) will result in amplification of any additional modification that was made to the data *after* that degradation took place, regardless of whether this modification is in the form of detector or photon noise, detector nonlinearities, or inhomogeneities in the response of any of the elements taking part in the acquisition of the data (such as filters or modulators). For this reason, it is advantageous to apply the instrumental effects to the synthetic data as a part of the relevant physics, rather than attempting to remove them from the data itself. The most common instrumental degradation treated in this way so far is the degradation in the spectral dimension. Since the effect of this degradation on the spectra is frequently very severe, convolution of the synthesized spectra with the instrumental transmission profile before comparing to the observed data is a standard ingredient in almost any inversion code.

A rudimentary treatment of spatial degradation effects can be found in the form of a micro- and/or macro-turbulent broadening of the spectral line profiles in almost all inversion codes; however, a thorough treatment of degradations in the spatial domain has thus far not been employed. This is probably not due to the relative importance of spatial degradation as compared to spectral degradation, but more likely due to the limitations imposed by the available datasets. Recent advances in instrumentation (Hinode SP, IBIS, CRISP, etc), have produced datasets sampling both the spatial and the spectral dimensions with high enough resolution to open up the possibility of improved treatment.

Orozco Suárez et al. (2007) used the observed data from neighboring pixels to assemble a "local stray-light" spectrum, which contributes to the observed spectrum in a particular pixel of Hinode SP slit-scanned data. As acknowledged by the authors, this method is far from exact, but is simple to understand and easy to implement, and it fixes most of the fitting problems they encountered.

However, since the estimate of the "stray-light" contribution can only be calculated from the observed data, this method in effect approximates a deconvolution by subtracting a running average. Apart from being formally incorrect, this procedure has the notable disadvantage that it removes an estimate of the signal not originating in a particular pixel (local stray-light) from the observations, thereby substantially decreasing the signal-to-noise ratio of the data. In addition, the process of using the data to calculate a fit parameter (the stray-light fraction) makes the solution vulnerable to systematic errors (Asensio Ramos & Manso Sainz 2011).

In the following, we develop an integrated approach to the problem of inverting spatially and spectrally degraded data. To see how to do this, we must first analyze the problem in some more detail.

2.1. Image response

The collection of solar spectral imaging data is no different from any other imaging problem, in that an image of the solar surface is formed by a telescope through an aperture in a focal plane, where it is sampled by an instrument that can detect both the spatial and spectral variability of the incident radiation. As a

consequence of the finite aperture of any imaging device, the radiation originating in a point on the object is redistributed across a spatially extended region on the detector.

The implication of this for two neighboring points of the object is that their respective spectra have been averaged before they are detected in the focal plane. In the probable case that the spectra differ, the result is a mixed spectrum that may not even resemble the original spectra.

To describe the effect of the spatial smearing with a PSF $\phi(x, y, \lambda)$, we consider the data observed if only a single point source is present in (x_0, y_0):

$$d(x, y, \lambda) = i(x_0, y_0, \lambda)\phi(x - x_0, y - y_0, \lambda).$$

For a spectrum that is completely described by a set of parameters α_i we may then write for the response of the observed data to a change of the parameter α_i in location (x, y):

$$\frac{\partial d(x, y, \lambda)}{\partial \alpha_i} = \phi(x - x_0, y - y_0, \lambda)\frac{\partial I(x_0, y_0, \lambda)}{\partial \alpha_i} \tag{1}$$

since the instrumental PSF does not depend on the α_i. We recognize the derivative on the RHS as the response function of the modeled spectrum used in most inversion methods employing a "greedy" algorithm (see for instance Cormen et al. 2001, for a description) to minimize a merit function (Landi Degl'Innocenti & Landi Degl'Innocenti 1977). The above expression gives us the true response of the data, given the undegraded response of the spectrum. We use this below to formulate the correct response matrix for a 2-D field of point sources, degraded by a known PSF.

Although not strictly necessary, in the following we assume that the spatial part of the PSF does not depend on the wavelength so that we can write

$$\phi(x, y, \lambda) = \varphi(x, y)\psi(\lambda).$$

There are many situations where this assumption is not valid, but it allows us to carry out the convolution over the wavelength before considering the spatial coupling, which makes the implementation somewhat simpler. We intend to relax this restriction in future work.

2.2. Levenberg-Marquardt minimization

The Levenberg-Marquardt minimization procedure (Levenberg 1944; Marquardt 1968) is used extensively in many areas where the minimum of a nonlinear function needs to be located. It uses an adjustable mixture of a fully linearized solution and a steepest descent gradient search to locate the global minimum of a function of an arbitrary number of variables. It's robust convergence properties make it the method of choice for many atmospheric inversion codes, and it is also at the core of the SPINOR code.

The method is built on a linearization of the fit-quantity (the spectrum) around a "best guess" set of fit parameters (the atmospheric parameterization parameters). It requires computing the derivative of each synthesized data point y_i with respect to each fit parameter α_i, each of which composes one element of the so-called Jacobian matrix $\mathbf{J}$.

Although the calculation of $\mathbf{J}$ is generally possible only by using a finite-difference approximation with respect to the fit quantity, the specific form of radiative transfer equation (RTE) allows us to compute them by moving the derivatives inside the formal solution integral (Landi Degl'Innocenti & Landi Degl'Innocenti 1977; Rees et al. 1989; Ruiz Cobo & del Toro Iniesta 1992) and evaluating the RTE using the derivatives of the opacity and emissivity with respect to the atmospheric quantities instead of the opacity and emissivity themselves. The resulting response functions (RFs) can be computed at a relatively modest additional cost alongside the emergent spectrum and correspond to the partial derivatives of the spectrum with respect to the fit quantities, required by the Levenberg-Marquardt minimization procedure.

Assuming the response functions are known, the solution α = must describe the observed data $\mathbf{d}$

$$\mathbf{J}\alpha = \mathbf{d}.$$

For an approximate solution α'

$$\mathbf{J}\alpha' = \mathbf{d}' \neq \mathbf{d}.$$

so that by using the linearity of $\mathbf{J}$ we may write

$$\mathbf{J}(\alpha - \alpha') = \mathbf{J}\nu = \mathbf{d} - \mathbf{d}' = \beta, \tag{2}$$

where ν is the linear correction for the set of atmospheric parameterization values α and β is the difference between the observed data and the spectrum synthesized using α. Since we generally have more data points then fit parameters, our desire to invert $\mathbf{J}$ directly to obtain ν is frustrated by $\mathbf{J}$ not being square. The solution is to multiply both sides by the transpose matrix $\mathbf{J}^T$, thus creating the pseudo-inverse problem

$$\mathbf{J}^T\mathbf{J}\nu = \mathbf{J}^T\beta, \tag{3}$$

where multiplying β by $\mathbf{J}^T$ can be seen as converting the overdetermined but exact problem into a least square fitting problem. Although this linear system gives us the least-square solution to the linearized problem, unfortunately, the original inversion problem is usually far from linear, so that direct application of this solution may be anything but an improvement. An additional damping constant λ_M is therefore added to (3) to control the behavior. The result is the familiar Levenberg-Marquardt form

$$(\mathbf{J}^T\mathbf{J} + \lambda_M \mathbf{diag}[\mathbf{J}^T\mathbf{J}])\nu = \delta, \tag{4}$$

with $\delta = \mathbf{J}^T\beta$ and λ_M the Marquardt damping constant. The matrix on the RHS can now be inverted to solve for the correction vector ν. It is easy to see that for increasing values of the damping constant λ_M, the fit matrix becomes more and more diagonal, while the corrections become smaller and smaller, thus approaching the steepest descent minimum search, which is known to be very slow and very vulnerable to finding a local minimum that is not the global minimum, but is guaranteed to lead to a decrease in the value of the merit function.

The convergence strategy is now clearly to keep the damping parameter as small as possible, in order to avoid the slow convergence properties of the steepest descent method, but large enough not to overstep the minimum so far as to increase the merit function. Ideally, the initial steps should only be concerned with the large-scale behavior of the minimization function; small-scale details should only come into focus once the dominant features have already been fairly well fitted. Unfortunately, as with all blind search algorithms, no guarantee exists that the minimum that is located is the lowest minimum.

2.3. Spatial coupling

Now that we understand the way in which the Levenberg-Marquardt method optimizes the merit function, we are ready to consider the spatial coupling between the pixels in an actual dataset.

There are at least two important ways in which the spectra in adjacent pixels depend directly on each other: horizontal radiative transfer effects and instrumental scattering and diffraction effects. The first involves the considerable task of solving the RTE in non-LTE and in full 3-D and computing the response functions for all image elements and all fit parameters in these image elements. Although it is not clear a priori that this effect does not have to be consider in general, we assume here that there are at least some photospheric lines for which the assumption of LTE is approximately valid so this effect can be neglected. The second way is easily seen to constitute a much simpler problem than the first, since the local derivative of the spectrum to the atmospheric fit parameters is simply spread to the neighboring pixels according to the PSF.

Taking a closer look at the Jacobian matrix of a spectro-polarimetric image makes it clear that for spatially uncoupled inversions, without a PSF, the Jacobian is simply an assembly of uncoupled sub-Jacobians:

$$\mathbf{J} = \begin{pmatrix} \mathbf{J}_{1,1} & & \cdots & & 0 \\ & \mathbf{J}_{1,2} & & & \\ \vdots & & \ddots & & \vdots \\ & & & \mathbf{J}_{n,m-1} & \\ 0 & & \cdots & & \mathbf{J}_{n,m} \end{pmatrix} \tag{5}$$

so that we simply have

$$\mathbf{J}^T = \begin{pmatrix} \mathbf{J}^T_{1,1} & & \cdots & & 0 \\ & \mathbf{J}^T_{1,2} & & & \\ \vdots & & \ddots & & \vdots \\ & & & \mathbf{J}^T_{n,m-1} & \\ 0 & & \cdots & & \mathbf{J}^T_{n,m} \end{pmatrix} \tag{6}$$

and the required fit matrix $\mathbf{J}^T\mathbf{J}$ clearly remains block-diagonal and can be solved efficiently block-by-block.

For spatially coupled inversions with a *uniform* PSF $\varphi(x, y)$ that only depends on the spatial coordinates, the Jacobian can be derived by making use of (1) and can still be written in a blocked form:

$$\mathbf{J} = \begin{pmatrix} \varphi_{0,0}\mathbf{J}_{1,1} & \varphi_{0,-1}\mathbf{J}_{1,2} & \cdots & \varphi_{1-n,2-m}\mathbf{J}_{n,m-1} & \varphi_{1-n,1-m}\mathbf{J}_{n,m} \\ \varphi_{0,1}\mathbf{J}_{1,1} & \varphi_{0,0}\mathbf{J}_{1,2} & & \varphi_{1-n,3-m}\mathbf{J}_{n,m-1} & \varphi_{1-n,2-m}\mathbf{J}_{n,m} \\ \vdots & & \ddots & & \vdots \\ \varphi_{n-1,m-2}\mathbf{J}_{1,1} & \varphi_{n-1,m-3}\mathbf{J}_{1,2} & & \varphi_{0,0}\mathbf{J}_{n,m-1} & \varphi_{0,-1}\mathbf{J}_{n,m} \\ \varphi_{n-1,m-1}\mathbf{J}_{1,1} & \varphi_{n-1,m-2}\mathbf{J}_{1,2} & \cdots & \varphi_{0,1}\mathbf{J}_{n,m-1} & \varphi_{0,0}\mathbf{J}_{n,m} \end{pmatrix},$$

but is clearly no longer block-diagonal. In practice, the extent over which φ does not vanish is much smaller than the computational domain, resulting in a relatively sparse system. The transpose is readily written down as

$$\mathbf{J}^T = \begin{pmatrix} \varphi_{0,0}\mathbf{J}^T_{1,1} & \varphi_{0,1}\mathbf{J}^T_{1,1} & \cdots & \varphi_{n-1,m-2}\mathbf{J}^T_{1,1} & \varphi_{n-1,m-1}\mathbf{J}^T_{1,1} \\ \varphi_{0,-1}\mathbf{J}^T_{1,2} & \varphi_{0,0}\mathbf{J}^T_{1,2} & & \varphi_{n-1,m-3}\mathbf{J}^T_{1,2} & \varphi_{n-1,m-2}\mathbf{J}^T_{1,2} \\ \vdots & & \ddots & & \vdots \\ \varphi_{1-n,2-m}\mathbf{J}^T_{n,m-1} & \varphi_{1-n,3-m}\mathbf{J}^T_{n,m-1} & & \varphi_{0,0}\mathbf{J}^T_{n,m-1} & \varphi_{0,1}\mathbf{J}^T_{n,m-1} \\ \varphi_{1-n,1-m}\mathbf{J}^T_{n,m} & \varphi_{1-n,2-m}\mathbf{J}^T_{n,m} & \cdots & \varphi_{0,-1}\mathbf{J}^T_{n,m} & \varphi_{0,0}\mathbf{J}^T_{n,m} \end{pmatrix},$$

which, using $\varphi * \varphi = Y$, allows us to write the desired matrix as

$$\mathbf{J}^T\mathbf{J} = \begin{pmatrix} Y_{0,0}\mathbf{J}^T_{1,1}\mathbf{J}_{1,1} & \cdots & Y_{1-n,1-m}\mathbf{J}^T_{1,1}\mathbf{J}_{n,m} \\ Y_{0,1}\mathbf{J}^T_{1,2}\mathbf{J}_{1,1} & & Y_{1-n,2-m}\mathbf{J}^T_{1,2}\mathbf{J}_{n,m} \\ \vdots & \ddots & \vdots \\ Y_{n-1,m-2}\mathbf{J}^T_{n,m-1}\mathbf{J}_{1,1} & & Y_{0,-1}\mathbf{J}^T_{n,m-1}\mathbf{J}_{n,m} \\ Y_{n-1,m-1}\mathbf{J}^T_{n,m}\mathbf{J}_{1,1} & \cdots & Y_{0,0}\mathbf{J}^T_{n,m}\mathbf{J}_{n,m} \end{pmatrix} \tag{7}$$

which contains the matrix product of all $\mathbf{J}^T_{i,j}$ with all $\mathbf{J}_{x,y}$, weighted with the autocorrelation function of the PSF centered on the location of the $\mathbf{J}_{x,y}$.

If the PSF is not uniform, it is straightforward to replace the coefficients $\varphi_{i,j}$ with appropriate vectors containing the elements of the appropriate PSF. Although the computational overhead of this is probably not prohibitive, we leave this to be developed in a future paper.

The Herculean effort of directly inverting this linear system, containing $(n_a \times n_s \times n_x \times n_y)^2$ elements, must not only be considered well out of reach of any computing resource likely to be available in the near future, if it were accomplished by some means, it is very likely that the large number of operations needed for computing the solution vector would lead to acumulation of unacceptable numerical errors, since the inverse matrix would likely not be sparse.

However, since the forward problem *is* sparse, we are able to accurately evaluate

$$(\mathbf{J}^T\mathbf{J} - \lambda\,\mathbf{diag}[\mathbf{J}^T\mathbf{J}])\mathbf{x}_n = \mathbf{A}\mathbf{x}_\mathbf{n} = \delta_n \tag{8}$$

and compute the defect $\mathbf{d}_n$ in the current estimate $\mathbf{x}_n$ of the correction vector ν

$$\mathbf{d}_n = \delta - \delta_n \tag{9}$$

in a reasonable amount of time. We now approximate $\mathbf{A}$ with a sufficiently accurate approximation, $\mathbf{A}^*$, which is easy to invert but captures the essence of $\mathbf{A}$. The solution of

$$\mathbf{A}^*\zeta_\mathbf{n} = \mathbf{d}_\mathbf{n} \tag{10}$$

can now be used to improve the estimate

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \zeta_n,$$

repeatedly, until $|\,\mathbf{d}_n\,|^2 < \epsilon$, with ϵ a suitably small, positive number. This process converges rapidly if an accurate approximation of $\mathbf{A}$ is used.

2.4. Approximate operator

In the process outlined above, an approximation to the full matrix $\mathbf{A} = \mathbf{J}^T\mathbf{J} - \lambda_M\,\mathbf{diag}[\mathbf{J}^T\mathbf{J}]$ is required that captures the basic properties of the full operator but is both cheap to invert and has a sparse inverse. The block diagonal approximation is an obvious candidate, because it works well for a compact PSF, but requires a large number of iterations to converge when a more extended PSF is used.

An extended-block variant of this method was also tried, where a spatially connected region was selected and inverted explicitly, as illustrated in Figure 1. This method generally improves the convergence characteristics considerably, but comes at the increased cost of inverting a larger part of $\mathbf{A}$ explicitly. Apart from the increased cost of the inversion, this approach suffers from the limitation that when the size of the blocks becomes

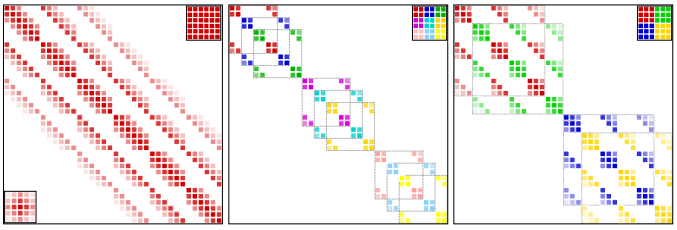

Fig. 1. Full matrix for a grid of 6×6 pixels (**left**), approximated by 9 blocks of 2×2 pixels (**middle**) and 4 blocks of 3×3 pixels (**right**). The cost increases rapidly with the number of pixels N, as the inversion of the resulting matrix scales $\propto N^3$. In the upper right corner, the spatial distribution of the pixels is indicated by the color of the block they belong to. The full matrix has in addition the cross-correlation function, **Y**, inserted in the lower left corner.

too large, the numerical errors in the inverted block matrix can become large enough for the iterative approximation to fail to converge completely. It is therefore important to find the optimum balance between the cost of explicit inversion and the cost of applying a larger number of iterations, which generally depends on the problem.

Other variations on the extended block method can be made, when more nonlocal behavior needs to be captured adequately. The staggered block method from Figure 2 can be useful when

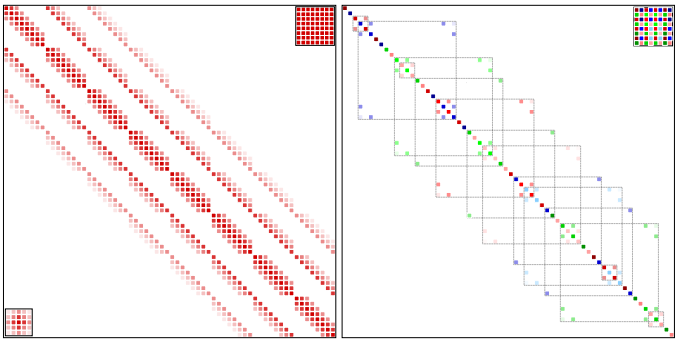

Fig. 2. Full matrix for a grid of 8×8 pixels (**left**), approximated using a staggered configuration of 2×2 pixel blocks (**right**). The approximated matrix has more than 9 blocks, due to boundary effects, compared with only 9 above, but many of the blocks are smaller. The insets are as in Figure 1

interaction over long distances is important, such as for significantly oversampled data. Intermittent application of several methods is possible, but the need to recompute the inverse of the approximate matrix generally outweighs the possible gain in convergence.

Computational optimization using explicit storage of all nonzero elements of $(\mathbf{J}^T\mathbf{J}-\lambda_M\,\mathbf{diag}[\mathbf{J}^T\mathbf{J}])$ and the approximation of the inverse, in combination with multithreading of the iterative solver, reduces the additional time overhead of this method to only ≈30% of the cost of synthesizing the Stokes profiles and response functions with the SPINOR code, for critically sampled data and with observational data containing ∼100 spectral points. This overhead is expected to decrease further when the amount of spectral points increases.

Depending on the approximate operator that is used, the convergence of the inversion of **A** typically requires on the order of 100 iterations to reach an error in the defect of 10%, as shown in Figure 3. This takes significantly less time than the computation of **A** and $\mathbf{A}^*$ and so does not significantly contribute to the overall execution time, as long as the inverted submatrices can be stored in RAM.

Unfortunately, the need to update the correction vector and recompute the defect, requires knowledge of the current estimate across the extent of the **Y** functions, resulting in the need for a significant amount of communication when a distributed parallelization strategy is desired. For this reason, at present, only an efficient multithreaded algorithm exists, requiring a shared memory environment.

In some cases, additional shortcuts can be taken to reduce the work requirements. Such shortcuts, however, do not always result in unconditional convergence, as is the case for the complete treatment as outlined above. One particularly useful approximation in the case of a PSF with very extended wings is to calculate the defect β explicitly using the complete PSF, while calculating the Levenberg-Marquardt correction ν using only the central part of the PSF. This results in good convergence when the value of the PSF integrated over the extended part of the PSF does not exceed the value of the PSF integrated across the central part. If this condition is not met, an exponentially growing correction is obtained and the inversion process will diverge.

2.5. Inversion strategy

Although we can evaluate the global problem, it is not self-evident that this is the optimal way to proceed. If we consider the pixel-by-pixel approach as completely uncoupled and a very extended PSF as a global inversion problem, it is clear that in general all levels of coupling will be encountered. For the calculations presented in this paper, only the global problem has been considered; however, it should be kept in mind that the global approach may not be the optimal choice in all cases.

The value of the merit function, χ^2, is plotted in Figure 3 for the inversion of an MHD simulation as a function of the iteration

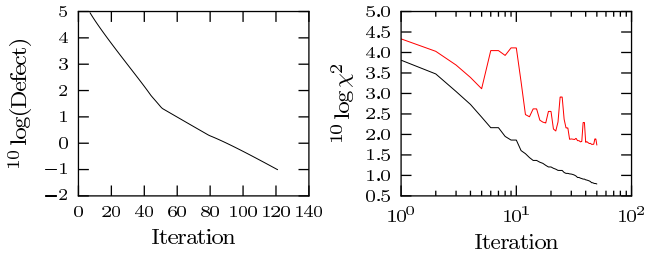

Fig. 3. Convergence properties of the iterative step determination and the global inversion **left**: error in the defect as a function of iteration number, **right**: average (black) and maximum (red) χ^2 as a function of iteration number.

number. The convergence rate can be strongly influenced by employing additional smoothing at regular intervals, which can effectively push the solution out of local minima by damping out large local differences that typically result from the cancellation of errors by other errors with the opposite sign.

A returning issue when converging oversampled data is the development of a large number of local minima, frequently involving spatial frequencies in the solution exceeding the diffraction limit. One of the more effective ways to reduce the probability of getting stuck in one of these minima is to initially limit the corrections applied to the solution to the lower frequencies and to relax this limitation as the inversion progresses.

Another effective way to avoid a slow and cumbersome descent to the global minimum is to introduce a significant amount of noise in the data in the first stages of the inversion. The noise effectively hides the less significant spectral features from the view of the inversion algorithm, leading to a rapid first descent phase. After the minimum has been approached with sufficient accuracy, the noise is gradually reduced, allowing for the details to enter the solution. This method is fairly effective in avoiding getting captured by outlying local minima, although it is not very effective in avoiding false minima near the global minimum.

3. Simulations

To evaluate and optimize the performance of the spatially coupled inversion problem, it is necessary to know what the value of a particular atmospheric parameter really is, something that can only be done in the context of simulated atmospheres. The atmosphere chosen for the tests below is a 288x288x100 point MURaM MHD simulation (Vögler et al. 2005) of the quiet Sun with a mean field of around 100G and a spatial resolution of 20.8km (0.029”). The grid size allows for convenient binning in a number of coarser grid sizes, which we make use of in the tests carried out below.

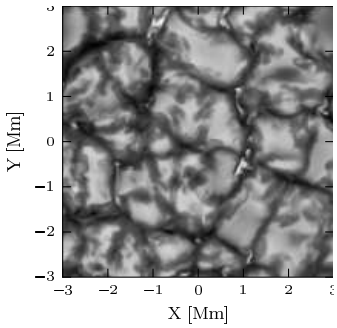

Fig. 5. The continuum intensity of the employed MURaM quiet Sun simulation.

The verification procedure necessarily remains strictly limited to validation and verification of the method under inherently ideal circumstances, since the forward and the backward problems use the same physics, a situation that is not likely to occur in the case of real observations. The effects of errors in the physics will be considered inherent to all inversions and are left for a later discussion.

The starting point of our calculations is an MHD cube, pictured in Figure 5, that provides the relevant atmospheric quantities (temperature, pressure, velocity and magnetic field) as a function of three spatial coordinates. An unambiguous comparison of such a fully stratified atmosphere with an inverted simplified atmosphere is a complicated task by itself and one that is not unique to the coupled inversion problem.

To steer clear of a complicated discussion on this topic, the atmosphere was first used to synthesize the emergent spectra, which were then inverted using simple three node atmospheres, with the nodes at the same heights as they will be for the inversions. The resulting atmospheres are only approximations to the original ones, but they have the advantage that they can be represented exactly by the inversion code.

The fitted atmospheres are then used to generate a second data cube of emergent spectra, with the notable difference that these spectra can be reproduced exactly by the simplified three-node atmospheres used by the inversion code, so that it should be possible to invert the spectra perfectly to the input atmospheres used to generate them, provided the mapping of the atmosphere to the emergent spectra is unique.

3.1. Ideal atmospheres

As motivated and explained above, we attempt to validate the method under the simplest possible circumstances, where the profiles are generated from a known solution that can be reproduced exactly by the inversion code. To achieve this, the three-node atmospheres, fitted to the synthesized spectra as described above, were first binned to the resolution of the artificial dataset we want to generate, *before* the spectra to be degraded and inverted are synthesized from them.

If the inversion problem is well-posed, this simplification should allow the inversion code to recover the binned atmosphere with arbitrary precision in the absence of spectral and spatial degradation. Since this is generally not the case for all pixels in the field of view (a few spectra will always be difficult to invert), to evaluate the properties of the coupled inversion as independently of the properties of the inversion procedure as possible, we focus on the differences between the pixel-by-pixel inversion of the undegraded data and the coupled inversion of degraded data.

3.1.1. Pixel-by-pixel inversions

After a first verification that the atmosphere used to generate the spectra gives a perfect match to the artificial data, a pixel-by-pixel inversion was carried out of the whole FOV, to test the inherent invertibility of the atmospheres using the algorithm used by the inversion code. Interestingly, the results indicate that, although the code should be able to perfectly reproduce the spectra, it does not succeed in doing so for approximately 1% of the pixels in the FOV.

Repeated “nudges” of the atmospheric parameters in random directions in parameter space improve the overall success rate, but in a number of pixels, the solution is evidently difficult to locate using the minimization procedure employed by the code, so that significant differences between the artificial observation and the inverted spectrum remain. However, where the inversion code indicates a good fit, the atmospheric parameters are accurately retrieved, suggesting that the inversion problem is not underdetermined in this simplified situation.

The fraction of the pixels that does not produce an accurate inverse is approximately independent of the resolution, which is somewhat surprising, since although the atmospheres are binned to the resolution of the data, so that no unresolved structure can be present, the profiles tend to exhibit more extreme behavior at higher resolution. Figure 4 shows the retrieved inverted values as a function of the input values. The temperature and Doppler velocity appear to be very well-determined in all layers of the atmosphere, but the magnetic field strength and orientation show minor scatter around the diagonal, indicating that there are limits to the invertibility of the problem.

Figure 6 shows the recovered Stokes parameters as a function of the input Stokes parameter. The agreement is very good in all Stokes parameters, indicating that the minor scatter observed in the recovered magnetic field parameters is due to [near] degenerate behavior in these parameters, possibly owing to an inherently low response.

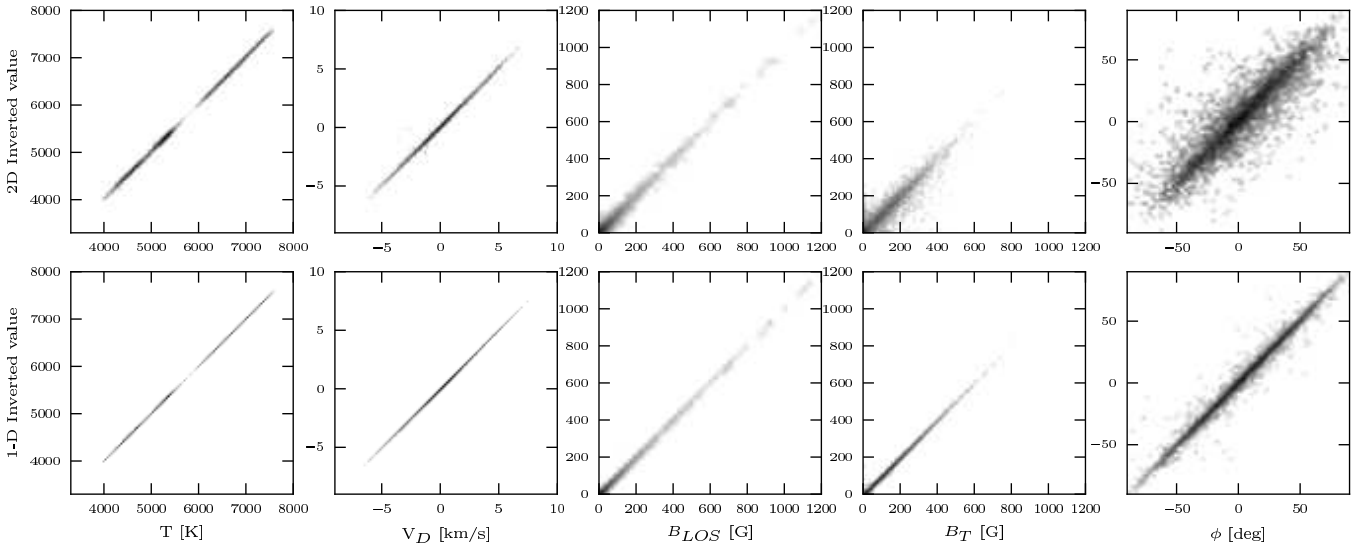

Fig. 4. 2-D histograms of inverted atmospheric parameters (vertical) as a function of input atmospheric parameters (horizontal), for a coupled inversion of simulated degraded 50cm telescope data (top row) and an uncoupled inversion of identical but *undegraded* simulated data (bottom row), for 5 atmospheric parameters across all heights in the atmosphere.

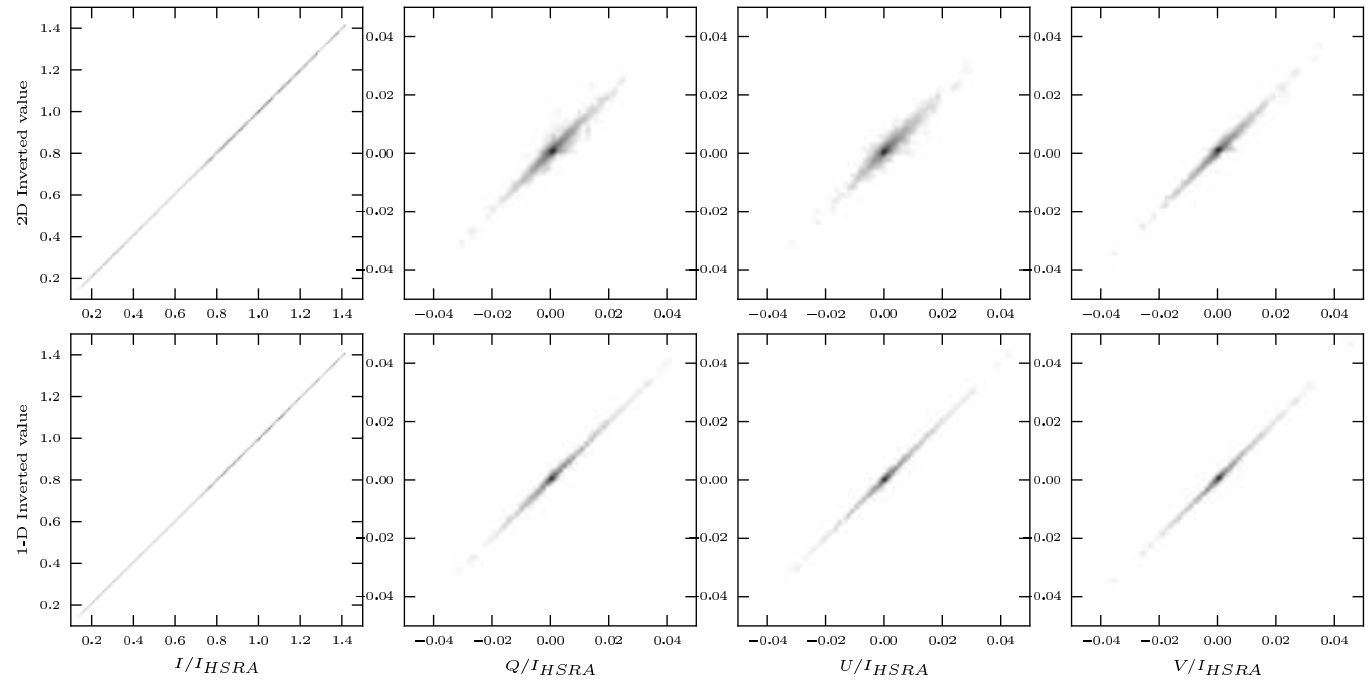

Fig. 6. 2-D histograms of inverted Stokes parameters (vertical) as a function of input Stokes parameters (horizontal), for a 2-D inversion of degraded 50cm telescope data (top row) and a 1-D inversion of identical but undegraded data (bottom row).

3.1.2. Spatially coupled inversions

In the spatially coupled case, we must choose a PSF to degrade the spectral data with. With a large volume of available potential data in mind, we choose the central part of the PSF of the SOT (Tsuneta et al. 2008) on board the Hinode spacecraft (Kosugi et al. 2007), as indicated in Figure 10, so we get a direct indication of the performance we can expect for such data.

The original resolution of the MHD cube, 0.029”/pixel, is approximately a factor 9 higher than the diffraction limit of 0.26” of the Hinode SOT at 6300Å, so that a binning factor of 5 would be the closest approximation of critical sampling. Since this is not an integer factor of 288, the cube was binned by a factor 6 to 48×48 pixels, corresponding to 0.174”/pixel, which is slightly undersampled and in reasonable agreement with the Hinode spectro-polarimeter (SP) (Lites et al. 2001) pixel size of 0.16”. In addition, a cube binned by a factor 3 to 96×96 pixels was produced to investigate the response of the algorithm to oversampled data with unresolved fine structure.

Compared to a pixel-by-pixel inversion, the spatially coupled inversions clearly require an increased amount of human intervention, since there are so many coupled fit parameters that there is a significantly increased probability of finding a local minimum that is not the global minimum. Identifying problem-

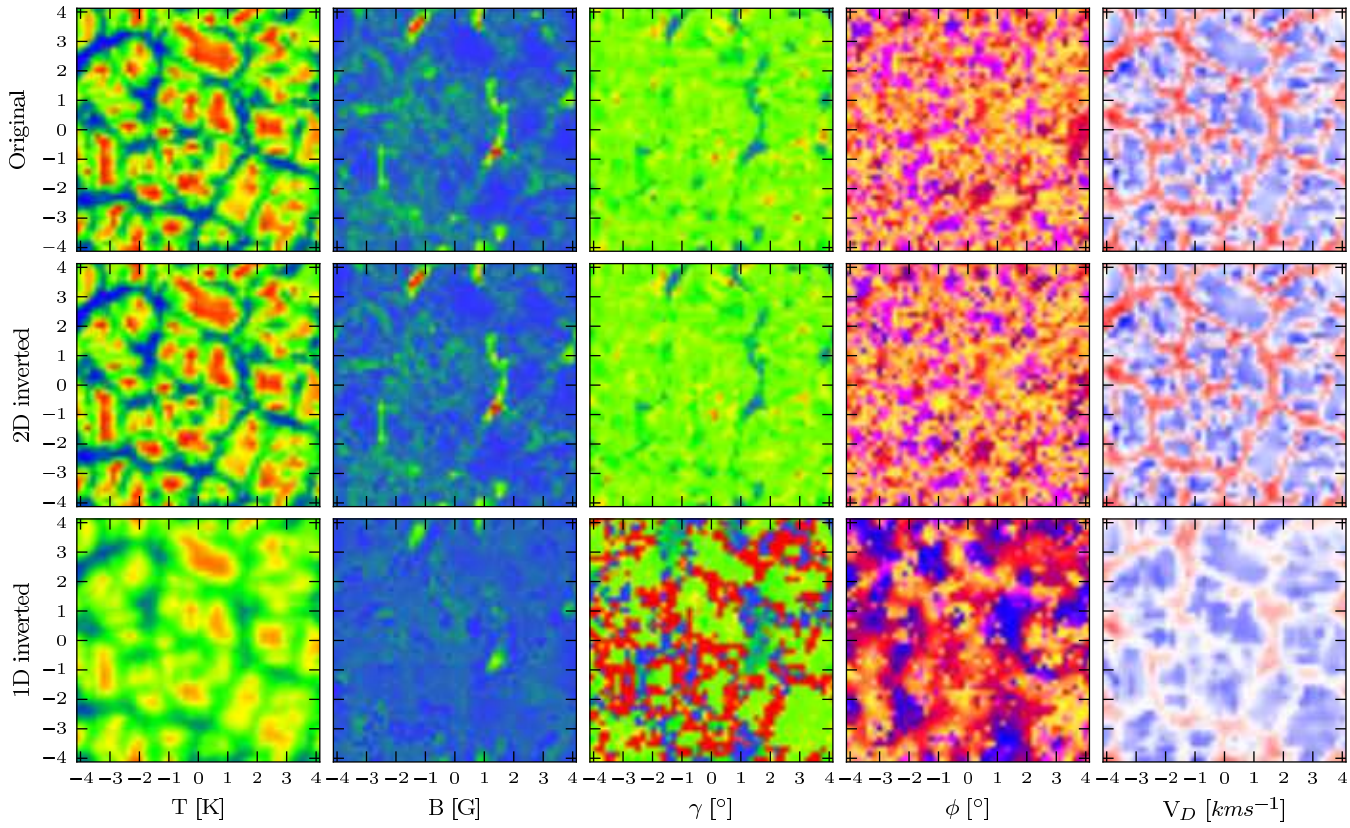

Fig. 7. Visual comparison of a selection of inverted parameters from spatially degraded spectra with the original data (top) for a spatially coupled inversion (middle) and a spatially uncoupled inversion (bottom). Shown are the parameter maps at the optical depths where they are most accurately determined. From left to right: Temperature at $\tau = 1$, magnetic field strength, inclination angle and azimuth angle at $\tau = 0.16$, and the line-of-sight velocity at $\tau = 1$. The color schemes are optimized for each quantity and have an identical but otherwise arbitrary scale for the original and inverted parameters. Tick marks indicate the spatial scale in arcseconds. Clearly visible is the difference in the accuracy of the retrieved parameters when the spatial smearing by the telescope is not taken into account.

atic pixels is difficult, because the only measure of the quality of the fit, the difference between the convolved profiles and the data, is the result of an ensemble of profiles, so that it is not possible to rate the contribution to the fit quality of one particular profile. A "nudge" of all fit parameters in a random direction is not necessarily helpful, since many pixels may have already converged and a random step in that solution may lead to improved convergence in one place, but not in another.

For slightly undersampled data, the method does not seem to have much more trouble recovering the solution than in the case of a 1-D inversion in the absence of any spatial degradation for the majority of the atmospheric parameters. Figure 4 shows the retrieved atmospheric quantities as a function of the input atmospheric quantities used to generate the undegraded spectra. Although the scatter around the diagonal is somewhat larger than for the pixel-by-pixel inversions of spatially undegraded data, when we keep in mind that the histograms in Figure 4 have a logarithmic gray scale, we may conclude that on a linear scale, the only parameters with significantly increased scatter are the horizontal components of the magnetic field.

This is also clearly visible in the inverted Stokes spectra shown in Figure 6, where the recovered Stokes Q and U values show the greatest deviation from their real value. Since the value of the merit function does not vanish for this solution, we can state with confidence that despite the human intervention, the inversion code found a local minimum that is not the global minimum, which is the most likely cause for the observed errors.

As a final illustration, Figure 7 shows the maps of a selection of the original parameters used to calculate the Stokes profiles, together with the inverted result. To emphasize the importance of correctly treating the spatial degradation, the uncoupled (1D) inversion result of the spatially degraded data is shown. Even though the parameters are shown at the optical depth where they are most accurately constrained by the data, the loss of accuracy caused by neglecting the spatial smearing of the data is clearly visible in the inverted results. We will explore the differences between the various inversion strategies in more detail in a future paper specifically on this topic.

3.2. The effect of noise

The effect of noise is one of the main problems to deal with when one is interested in the interpretation of real data. To investigate the effect of additive Gaussian noise on the spatially coupled inversions, a number of synthetic observations were produced with noise levels that are typical for data obtained with an efficient instrument and an integration time of several seconds.

The results of a repeat of the test from section 3.1.2, using data with varying levels of noise, is shown in Figure 8. The critical level for reliable recovery of the atmospheric parameters appears to be $10^{-4}\ I_c$, beyond which first the transversal (T) mag-

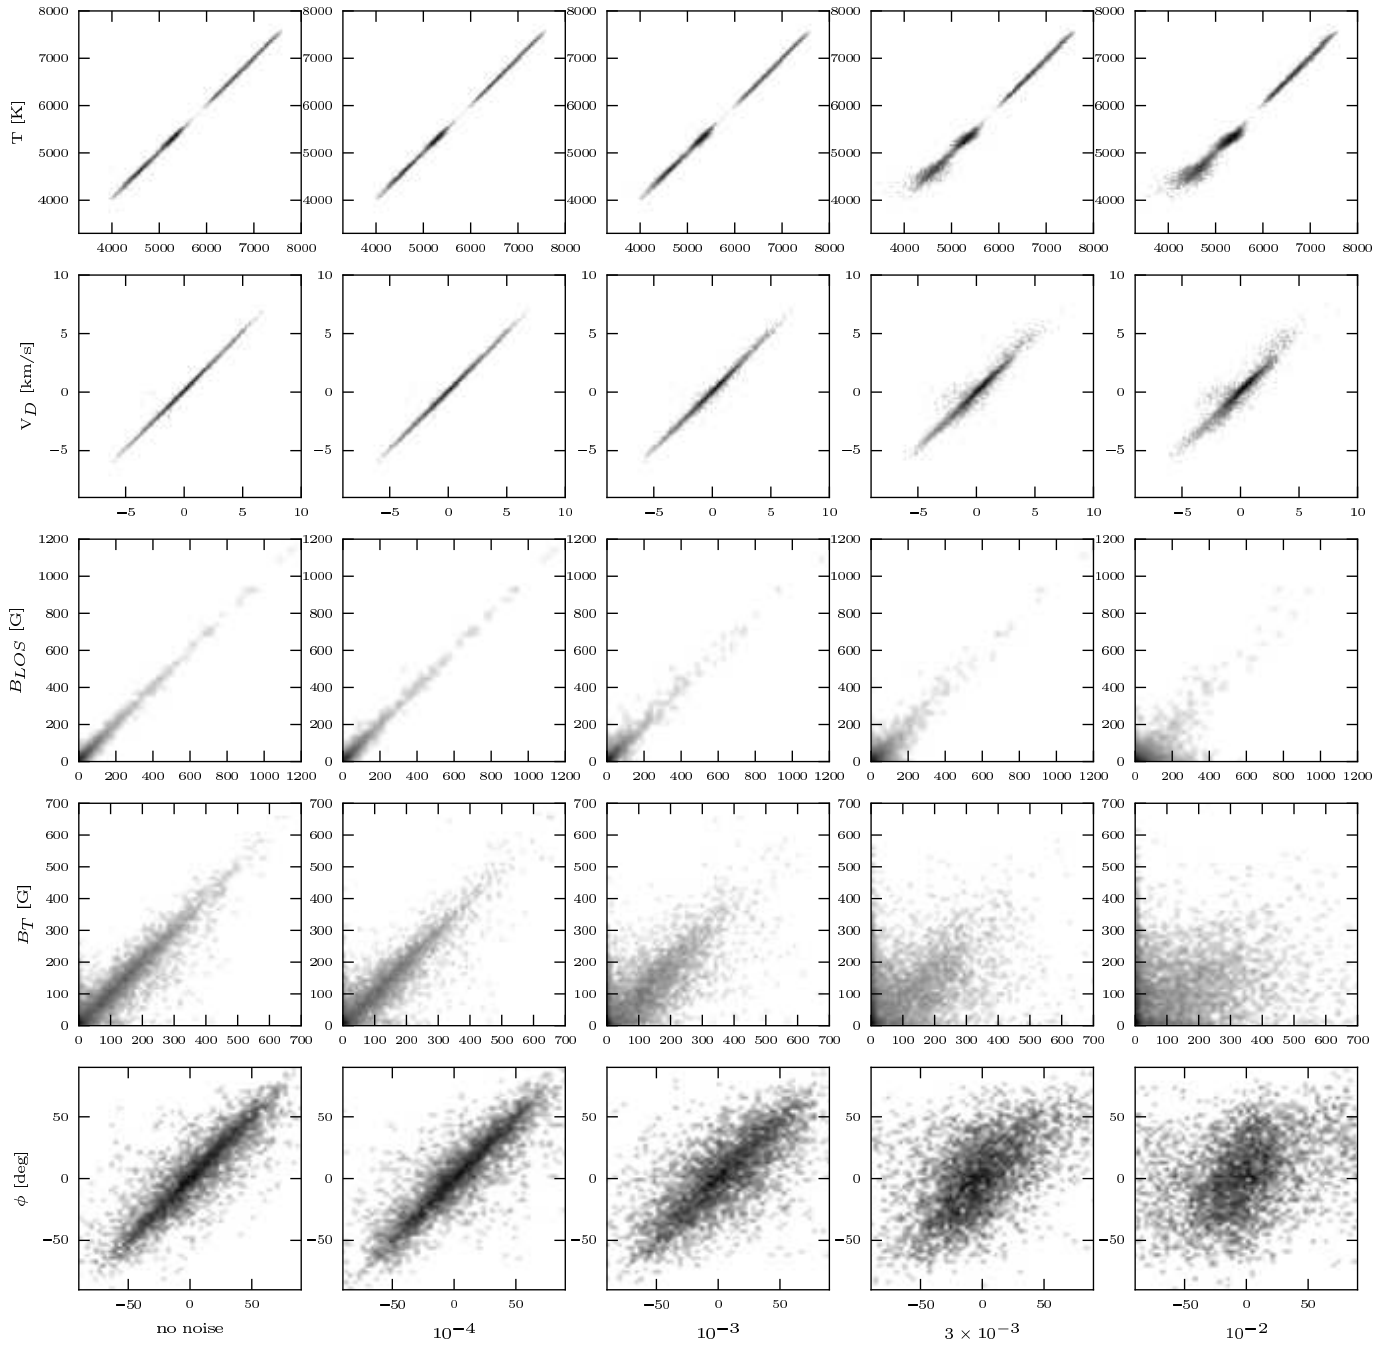

Fig. 8. Two-dimensional histograms of inverted atmospheric parameters (vertical) as a function of input atmospheric parameters (horizontal), for a 2-D inversion of degraded 50cm telescope data for noise levels of (left to right) 0, 10^{-4}, 10^{-3}, 3×10^{-3}, and 10^{-2} I_c, for 5 atmospheric parameters across all heights in the atmosphere.

netic field properties rapidly degrade, followed by the longitudinal magnetic field and finally the temperature and line-of-sight (LOS) velocity. At a noise level of 10^{-2} I_c, practically all atmospheric parameters show large errors, suggesting that a level well below that should always be achieved for any atmospheric property to be reliably recovered.

In particular, the recovered temperature and line-of-sight velocity appear to only be weakly affected by the noise. This is easily understood as these quantities may be inferred directly from the intensity, which is clearly the Stokes parameter least affected by the noise. In addition, the intensity in the continuum has a strong response to the temperature in the deepest layers, as well as a large number of spectral data points constraining that part of the profile.

The atmospheric quantities that only depend on the Stokes Q, U, and V parameters are significantly more sensitive to the noise, which is not only much more severe owing to the relatively low signal level of these data, but also to the negligible response of the majority of the (continuum) data points to these quantities. The final consequence of this is that for reliable recovery of the longitudinal component of the magnetic field in the quiet sun, a

signal-to-noise of $\sim 10^{-3}$ I_c is required, for the orientation and field strength it must be significantly better.

3.3. Unresolved structure

One of the main concerns in any inversion is the uniqueness of the solution, in other words, is the problem invertible at all? Naively one might expect the spatially coupled fit to have problems here, since only the sum of a number of spectra is constrained, not each individual spectrum. The difficulty the code has in locating the global minimum solution suggests, however, that this is actually not the case, so that if there are multiple optimum solutions, clearly none of them is found.

In an attempt to test what happens to the solution in cases where this uniqueness should break down, an inversion was done on oversampled simulated data, a case where conventional convolution is known to no longer provide one-to-one mapping. To this end, a data cube was simulated as before, but this time at a resolution of 0.085”/pixel, and convolved with the Hinode PSF sampled at 0.085”/pixel, so that it is possible for the solution to contain information that is not constrained since it is no longer present in the inverted observation after convolution with the PSF.

Not unexpectedly, the inversion does not find as good a solution as in the case of critically sampled data, even if just the value of the merit function is considered, but it shows no sign of instability or amplification of errors either. This suggests that the constraints provided by imposing a particular atmosphere and particular physics on the solution are able to stabilize the inversion sufficiently, even in oversampled cases.

Figure 9 shows the scatter of the inversion result as a function of the input atmosphere. Although the scatter is significantly increased over the critically sampled case, it is similar to the scatter of an appropriately spatially filtered atmosphere, shown in the bottom half of the same figure. The algorithm apparently does not recover much information beyond the diffraction limit, but instead appears to recover an atmosphere that has all frequencies beyond the diffraction limit removed from it.

3.4. Discussion

The simulation results show that the evaluation of the merit function in convolved data space allows for accurate discrimination between profiles in the undegraded image to spatial scales up to the diffraction limit of the PSF, if the instrumental properties are known. The undegraded Stokes profiles are recovered with a minor increase in the absolute error, a behavior closely followed by the inverted atmospheric parameters. Clearly, the quantities that depend primarily on the Stokes I profiles are recovered more accurately than those depending on Stokes Q, U, and V, even if only numerical noise is present.

Interestingly, the inversion algorithm does not converge to the absolute minimum, where the merit function was verified to vanish. The difficulty is that whereas a pixel-by-pixel inversion allows the identification of poorly converged pixels, this is not possible in the spatially coupled case.

The failure to find the absolute minimum within the space accessible by the algorithm suggests that if the solution is degenerate, the alternative degenerate solutions are as difficult to locate as the true solution, even in the case of oversampled data.

4. Application to Hinode observations

Since the spectral dimension is crucial for this method to work and the simplification introduced by the assumption of a constant PSF is advantageous in keeping the computations tractable, the spectro-polarimetric scans of the Fe I lines at 6301.5Å and 6302.5Å provided by the Hinode spectro-polarimeter were used to test the performance of the coupled inversion method on real solar data.

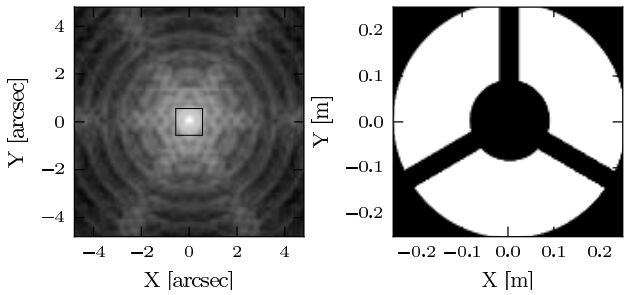

Fig. 10. left: The Hinode PSF at 0.16” per pixel resolution on a logarithmic gray scale from 10^{-6} (black) to 1 (white). The black box is the explicit part used in the spatially coupled inversions. **right:** Hinode SOT pupil function.

The significant scanning time of this instrument compared to the evolution time scale of some solar structures gives reason to be critical of the result, however, the slight under-sampling of the data keeps the evolution effects of the solar image relatively small on the critical spatial scale of the PSF. Nonetheless, it is important to remember that accurate reproduction of the spectra may not be achieved in some cases.

In addition to this obvious shortcoming, we are dealing with a real solar atmospheric structure, with, more likely than not, a complicated depth stratification, which will not be accurately represented by the simple atmosphere imposed by the inversion code. Furthermore, although presumably very accurate, the assumption of LTE probably does not hold exactly and the atomic and molecular line data are not known perfectly.

We nonetheless proceed by computing the PSF of the telescope. The PSF was computed from the Hinode pupil as specified in detail in the technical reference manual (Suematsu et al. 2008). As the focus position was found to differ from 0 by Danilovic et al. (2008), a 0.1 wave defocus was added, even though the focus position of this dataset was not accurately known. The PSF up to a radius of 4.8” is shown in Figure 10 on a logarithmic grayscale from 10^{-6} to 1, showing a complicated and extended interference pattern with significantly enhanced radial structures produced by the triangular spider.

To calculate a 2D inversion using the whole PSF, an unreasonably large amount of memory and computing time would be needed. Therefore, the shortcut described in section 2.4 was used, where the central part of the PSF explicitly used in the inversions is indicated by a black square in Figure 10. Since the wings of the PSF are taken into account explicitly when calculating the defect and the merit function, they are properly included in the converged result.

Since the results did not really require it, no straylight contamination was included in the inversion, which is not to say that there is none in the data. However, since the fit to the Stokes profile with the lowest intensity in the inverted data cube, which was found to be around $0.12 I_{c,HSRA}$ in the darkest part of the umbra,

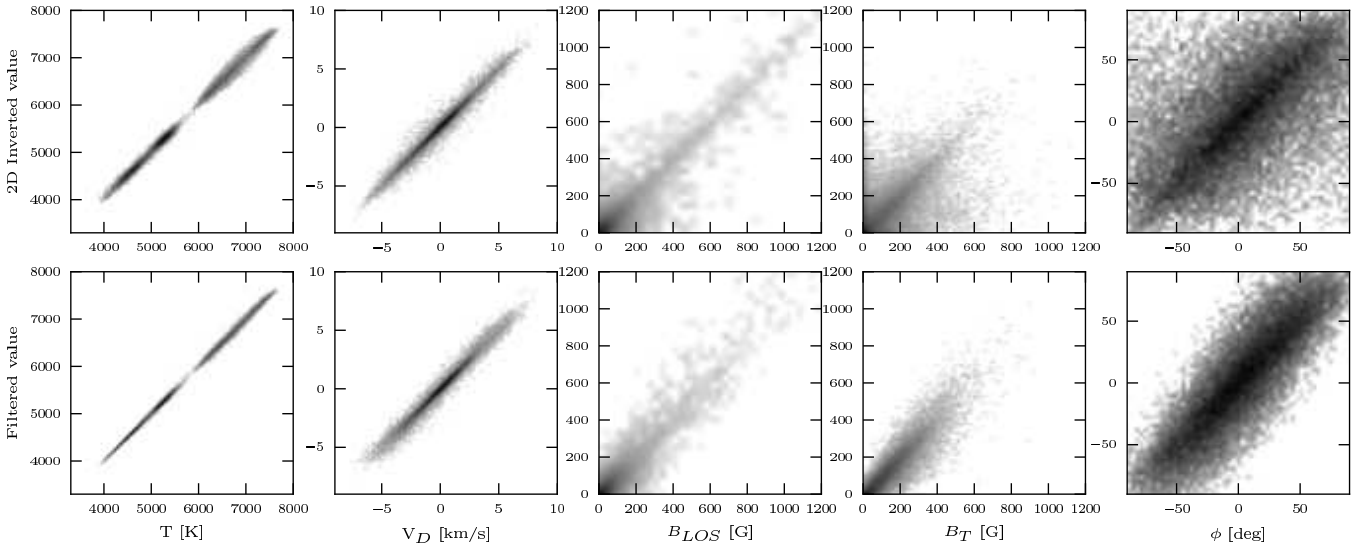

Fig. 9. Two-dimensional histograms of inverted atmospheric parameters (vertical) as a function of input atmospheric parameters (horizontal), for a 2-D inversion of degraded 50cm telescope data (top row) and a spatially filtered but otherwise undegraded atmosphere (bottom row), for 5 atmospheric parameters (horizontal) across all heights in the atmosphere.

was still accurate to within $\sim 0.1 I_{c,local}$, the amount of stray-light present that cannot be explained with the telescope PSF is unlikely to be more than around $0.01 I_{c,HSRA}$, although no effort was made to more precisely determine this number.

A scan of AR10933, recorded on 5 January 2007, 11:54-12:27 UT was selected, showing a sunspot umbra, penumbral filaments, and some solar granulation, very close to the center of the solar disk. The FOV was selected since it contains a wide variety of atmospheric conditions and fine structure that are likely to benefit from a spatially coupled inversion. As in the simulations from section 3, three nodes were placed at $^{10}\log\tau = -2.5, -0.9, 0.0$ to represent the solar atmosphere.

4.1. Inverted observations

Before the inversion results are discussed, it should be stated that the focus of this paper is a discussion of the coupled inversion method and not a discussion of the atmospheric structure that is the actual result of the inversion. Since discussing the latter would draw our attention away from the method, we leave a detailed analysis of these results to a separate paper and instead limit ourselves to a brief, qualitative assessment of the results combined with a comparison of the fitted undegraded profiles to the observed (degraded) data. More specifically, we produce images generated from the inverted profiles (inverted images) and compare them to images made from the observed data.

The inverted images at 6302.55Å of the selected active region are shown in Figure 11. The synthetic images generated from the inverted atmospheric structure (right side) compare very favorably with the original data (left side), in that they show much more contrast and a significant increase in the level of detail. This is even clearer in Figure 12, where the upper righthand corner of the umbra is shown enlarged for both the observed and the inverted data.

A selection of atmospheric parameters is shown in Figure 13, where the level of detail suffices to see the temperature drop in the penumbral dark cores at the highest node, the clear drop in magnetic field strength in the penumbral filaments, the large increase in inclination in the penumbral filaments, and clear signs

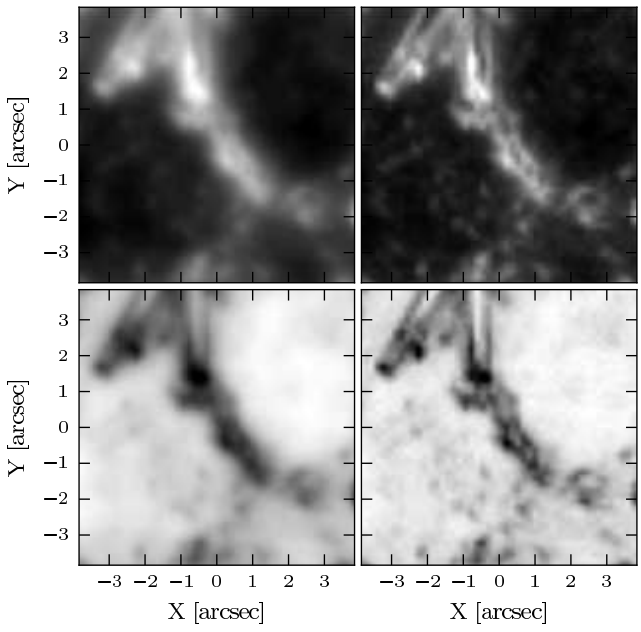

Fig. 12. Observed (left) and inverted (right) Stokes images in Stokes I (top) at and Stokes V (bottom) at 6302.55Å.

of downflow on both sides of the filament in the deepest layers, only seen before in datasets taken with larger solar telescopes (Joshi et al. 2011; Scharmer et al. 2011).

Despite this obvious improvement in the inverted parameters, the results also show several regions with obvious signs of oscillations on the scale of a single pixel. The oscillations show up in specific regions and cannot be removed by offering the code a smooth start condition. The pattern itself is confined to a certain region and does not spread beyond that, but its phase is not fixed. Several inversions starting from slightly different start conditions all produce a similar oscillatory pattern in the same region, but none of them is identical to any of the others. An av-

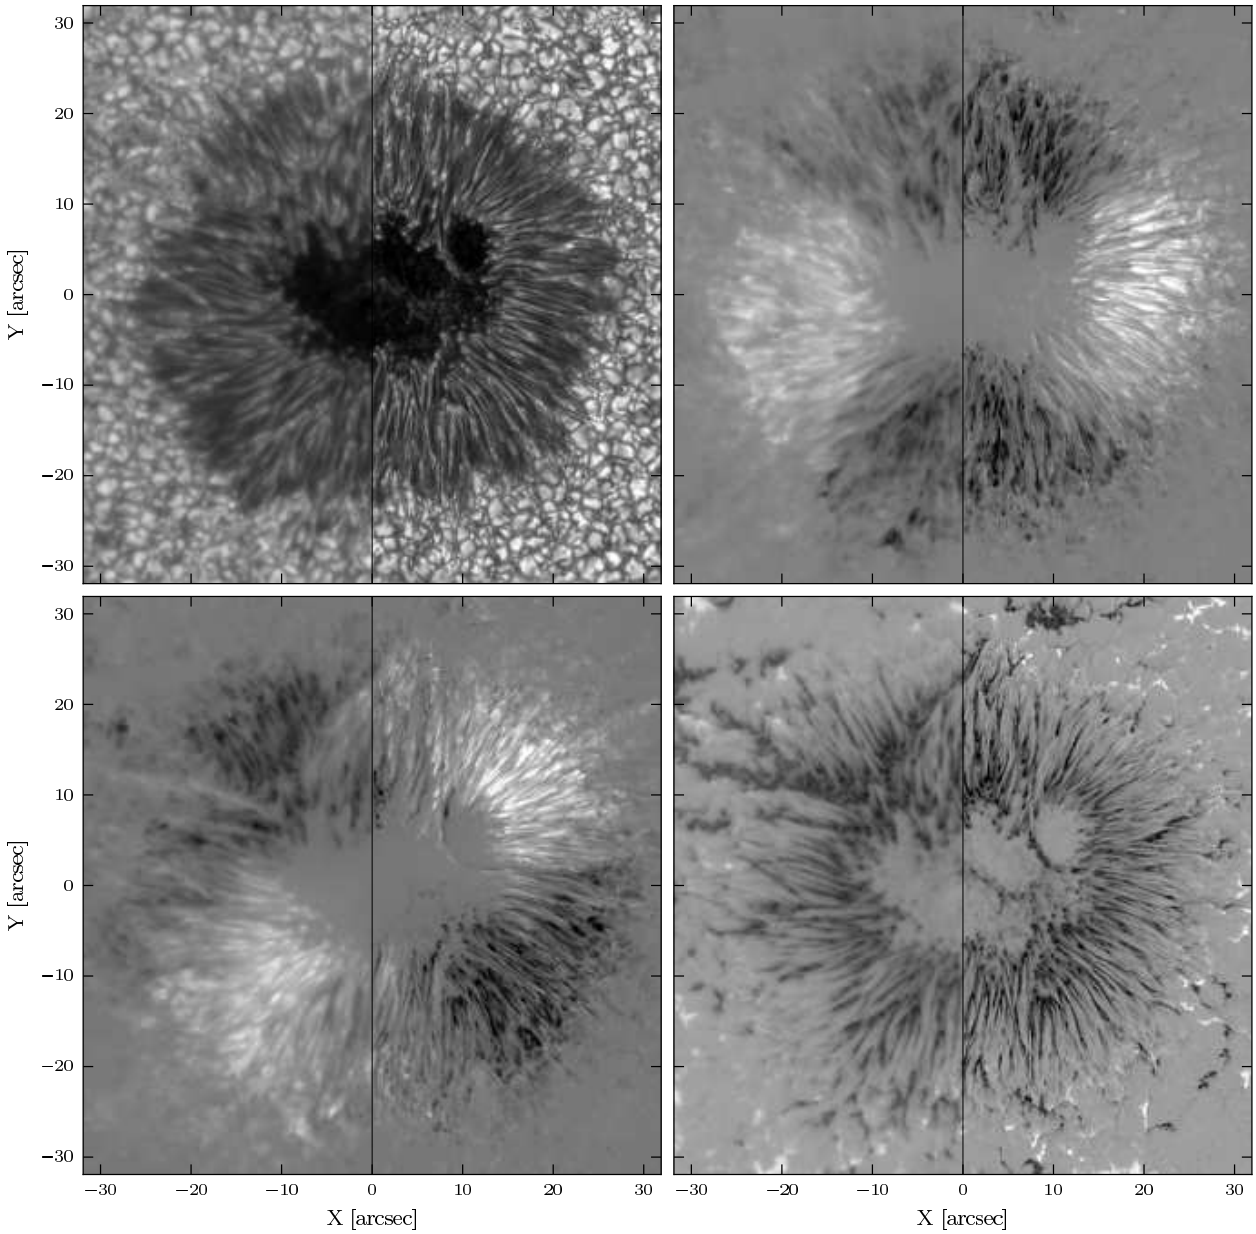

Fig. 11. Stokes images in the Fe I 6302.5Å line+ 57 mÅ, before (left half) and after (right half) inversion for the Stokes parameters I (top left), Q(top right), U(bottom left) and V(bottom right).

erage of several such solutions does not show any oscillations, but if an inversion is started from this averaged solution, it will develop new oscillations.

The regions in which these oscillations are found are typically fine structured and fast evolving, such as the outer penumbra. One example is shown in Figure 14, where the oscillations are clearly visible in the undegraded Stokes V and the LOS velocity in the penumbra, but are perfectly smooth in the granulation right next to it. The degraded Stokes V image shows that these oscillations average out to form a smooth image with only very weak oscillations.

The persistence of the oscillations in combination with the relative insensitivity of the inversion to their phase, but not their amplitude, suggest that we may not be dealing with just any degenerate solution here, but that the oscillations may actually be required to produce a spectral feature that cannot be produced by the atmospheric model. In particular, since the occurrence appears to be concentrated in the deep layers, the code must have difficulty accurately reproducing the wings of the line profiles. This may be due to calibration errors, but it is more likely caused by fine-scale structures that are not resolved by the telescope, possibly in combination with significant time evolution effects occurring during the spatial scan.

Apparently the inversion algorithm prefers to compensate for these errors by introducing what is in effect a microturbulent velocity structure, that is, by alternating the LOS velocity on the smallest possible scale. The averaging occurring in the convolution of the synthesized profiles then produces the broadened profiles that are observed. Interestingly, the removal of the microturbulent velocity as a fit parameter significantly increases the

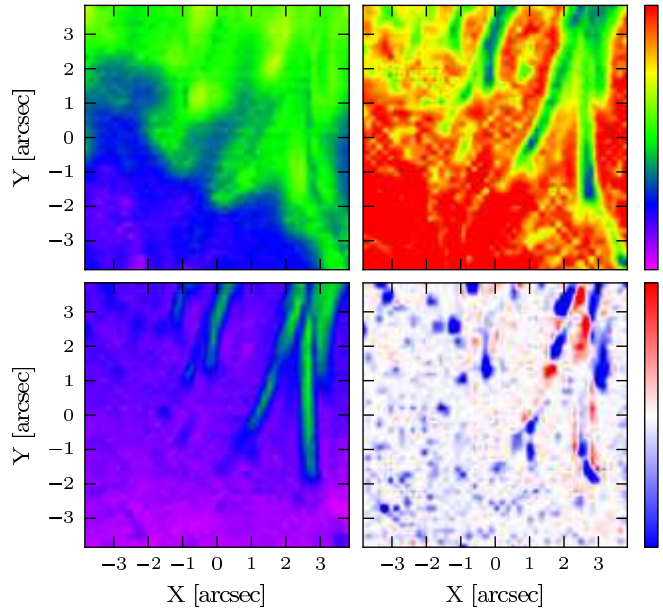

Fig. 13. A selection of inverted atmospheric parameters. From left to right, top to bottom: $T_{-2.5}$, $B_{-0.9}$, $\gamma_{-0.9}$, and $V_{D,0}$, with the subscript denoting the value of $^{10}\log\tau$. The color scale used for $T_{-2.5}$, $B_{-0.9}$, and $\gamma_{-0.9}$ is indicated on the top right, the one for $V_{D,0}$ on the bottom right. The ranges of the color scale are [3300-5200] K, [500-2500] G, [0-180] °, and [-3,3] $km\,s^{-1}$, respectively.

problem, suggesting that this parameter is used to fit some of the broadening caused by unresolved structure, but that significant differences remain that cannot be reproduced in this way.

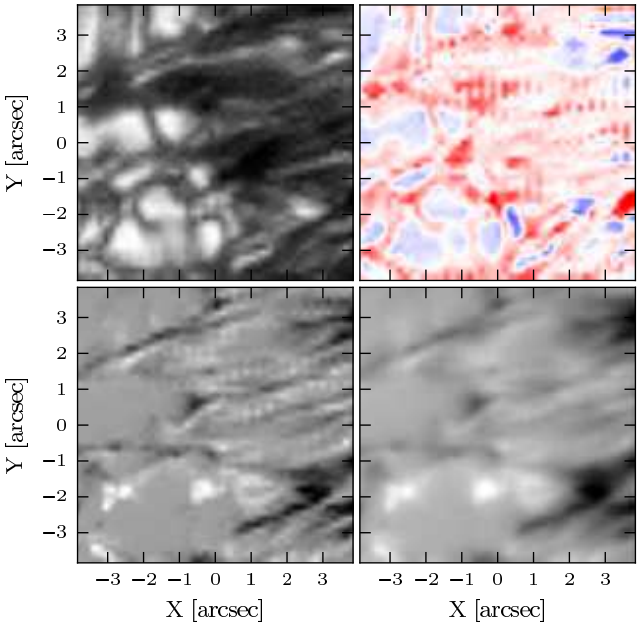

Fig. 14. Inverted Stokes images in Stokes I and V (left), alongside with the inverted LOS velocity at $\tau = 1$ and the degraded Stokes V. Clearly visible are the oscillations in the penumbra, which are invisible after degrading the image.

The temperature map in Figure 13 shows some horizontal stripes that can clearly not be associated with real solar features. They are parallel to the spectrograph scanning direction and can only be attributed to imperfections in the flatfields, made visible here by the level of detail to which the observed data are fitted. These artifacts are concentrated in the top layer of the inverted atmosphere, most likely owing to the small number of spectral points that are available to constrain this part of the atmosphere and the intrinsically lower intensity that these points have compared to the other parts of the spectrum.

Even a small calibration error in the width of the instrumental transmission profile will cause compensation effects in the solution, predominantly in the top layer. Indeed, a deliberate underestimate of the spectral transmission profile of the instrument introduces significant oscillations in the top layer of the atmosphere, as the algorithm attempts to broaden the line by alternating an atmospheric quantity.

It is clear that the algorithm is able to successfully fit the observed spectra, but it is able to do so with such high accuracy that it is essential to calibrate the data to a precision level that matches or exceeds that of the inversion result, which is significantly higher than the noise level of the data.

4.2. Deconvolved observations

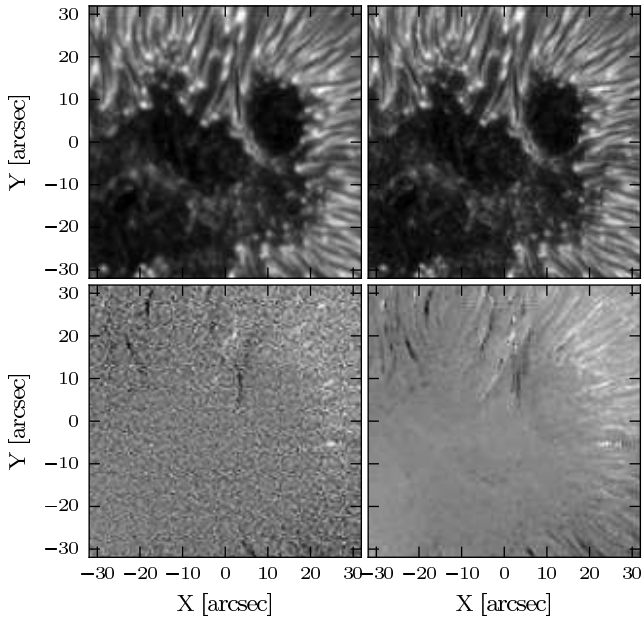

Fig. 15. Deconvolved (left) and inverted (right) Stokes images in Stokes I (top) at 6302.55Å (line wing) and Stokes U (bottom) at 6302.65Å (near continuum). Note the enhanced contrast in the high signal case (top) and the improved noise characteristics in the low signal case (bottom) of the inverted result as compared to the deconvolved result.

To evaluate the results obtained with the coupled inversions, the inverted undegraded profiles were compared to the data deconvolved using Richardson-Lucy deconvolution (Richardson 1972; Lucy 1974). Figure 15 shows two images, one deconvolved and the other one the result of the coupled inversion. The Stokes I images, at 6302.55Å, show a significant difference

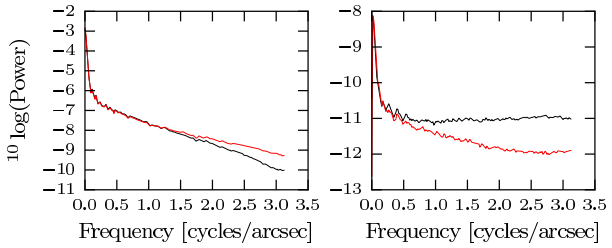

Fig. 16. Angular averaged power spectrum of the deconvolved (black) and inverted (red) solutions of Stokes I at 6302.55Å, approximately half way into the Fe I line at 6302.5Å. Clearly visible is the increased power of the inverted result in the high-frequency range of the power spectrum.

in contrast, as is confirmed by the azimuthally averaged power spectrum shown in Figure 16, indicating that the inverted image contains much more power at the high frequencies than the deconvolved result.

A low-signal image, on the other hand, such as the Stokes U image shown at the bottom of Figure 15, shows a power spectrum that has much more power than the inverted result, as can be seen in the righthand panel of Figure 16. Close inspection of the deconvolved image confirms that this can be attributed to a high level of noise in the image, not to any actual structure. The inverted result, on the other hand, still shows credible, coherent structures with nine orders of magnitude less power than the maximum power in the Stokes I image and does not contain any noise.

The inversion is able to produce high-frequency information when it is supported by the data, but suppresses it very effectively elsewhere. The likely reason for this lies in the fact that the inversion is based on the complete data cube, whereas the deconvolved result only contains information from a single wavelength. The effective noise level of the deconvolved result is therefore much higher than for the inverted result, specifically at very low signal levels, where the inverted result is constrained by the high-signal regions of the spectra, where noise is much less important.

5. Conclusions

We have developed a new method for inverting 2-D maps of spectro-polarimetric data that have been degraded spatially in a known way. The method uses the information contained in the spectral dimension and the known spatial degradation properties to constrain a parameterization of the atmosphere over the whole FOV simultaneously.

The method was formulated in a manner specific to an inversion code using a “greedy” optimization method, and although it is implemented for the SPINOR inversion code, it should not be considered as specific to this code. The method only changes the way in which the merit function is evaluated and modifies the response functions accordingly, implying that it should be easy to adapt to other inversion codes. The suitability of the specific implementation described in this paper is limited to the specific category of minimum search algorithms and cannot be used directly in Monte-Carlo or genetic codes without adaptations.

Tests carried out on simulated MHD cubes show that on a critically sampled grid, the code is able to retrieve atmospheric parameters by inverting profiles synthesized from them and degraded using a realistic PSF, with only a small increase in the error of the retrieved parameters as compared to a pixel-by-pixel inversion of the undegraded spectra. For oversampled data, the code is able to successfully recover the low-frequency components of the solution, but appears to have difficulty recovering information beyond the diffraction limit. At no time was unstable behavior resulting from oversampling the solution observed.

The method developed here is a specific case of a more generalized approach to data reduction, involving the application of calibration information to synthetic data (e.g. in the forward direction), instead of “correcting” the real data with the inverse calibration. Thanks to this forward application of the calibration information, the noise amplification typically associated with the correction of degraded data is strongly reduced, since the inverted calibration data are never applied to it. In addition, by combining all the available information to constrain a simplified parameterized atmosphere, all errors and limitations are accumulated in one place, resulting in a robust solution with considerably improved specifications over the input data.

Based on the results of section 4 we conclude that this method is suitable for application to spectral imaging data, where the assumption of a constant PSF can be made. There is no fundamental problem extending this method to the treatment of data where this assumption is not valid, which will be explored in a future paper.

Acknowledgements. The author wants to thank A. Lagg and S. Tiwari for the many helpful discussions that have resulted in the work presented here.

References

Asensio Ramos, A. & Manso Sainz, R. 2011, The Astrophysical Journal, 731, 125
Asensio Ramos, A., Manso Sainz, R., Martínez González, M. J., et al. 2012, The Astrophysical Journal, 748, 83
Asensio Ramos, A., Martínez González, M. J., & Rubiño-Martín, J. A. 2007, Astronomy & Astrophysics, 476, 959
Asensio Ramos, A., Trujillo Bueno, J., & Landi Degl’Innocenti, E. 2008, The Astrophysical Journal, 683, 542
Bernasconi, P. N. & Solanki, S. K. 1996, Solar Physics, 164, 277
Carroll, T. A., Kopf, M., & Strassmeier, K. G. 2008, Astronomy & Astrophysics, 488, 781
Cavallini, F. 2006, Solar Physics, 236, 415
Cormen, T. H., Stein, C., Rivest, R. L., & Leiserson, C. E. 2001, Introduction to Algorithms, 2nd edn. (McGraw-Hill Higher Education)
Danilovic, S., Gandorfer, A., Lagg, A., et al. 2008, Astronomy & Astrophysics, 484, L17
de Boer, C. R., Kneer, F., & Nesis, A. 1992, Astronomy & Astrophysics, 257, L4
Frutiger, C. 2000, Diss. ETH, 13896
Frutiger, C., Solanki, S. K., Fligge, M., & Bruls, J. H. M. J. 2000, Astronomy & Astrophysics, 358, 1109
Joshi, J., Pietarila, A., Hirzberger, J., et al. 2011, The Astrophysical Journal Letters, 734, L18
Knox, K. T. & Thompson, B. J. 1974, The Astrophysical Journal Letters, 193, L45
Kosugi, T., Matsuzaki, K., Sakao, T., et al. 2007, Solar Physics, 243, 3
Labeyrie, A. 1970, Astronomy & Astrophysics, 6, 85
Lagg, A., Ishikawa, R., Merenda, L., et al. 2009, in Astronomical Society of the Pacific Conference Series, Vol. 415, The Second Hinode Science Meeting: Beyond Discovery-Toward Understanding, ed. B. Lites, M. Cheung, T. Magara, J. Mariska, & K. Reeves, 327
Lagg, A., Woch, J., Krupp, N., & Solanki, S. K. 2004, Astronomy & Astrophysics, 414, 1109
Landi Degl’Innocenti, E. & Landi Degl’Innocenti, M. 1977, Astronomy & Astrophysics, 56, 111
Levenberg, K. 1944, Quarterly of Applied Mathematics, 2, 164
Lites, B. W., Elmore, D. F., Streander, K. V., et al. 2001, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4498, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. O. H. Siegmund, S. Fineschi, & M. A. Gummin, 73–83
Löfdahl, M. G. & Scharmer, G. B. 1994, A&AS, 107, 243
Lohmann, A. W., Weigelt, G., & Wirnitzer, B. 1983, Appl. Opt., 22, 4028
Lucy, L. B. 1974, AJ, 79, 745

Marquardt, D. 1968, SIAM Journal on Applied Mathematics, 11 (2), 431
Orozco Suárez, D., Bellot Rubio, L. R., Del Toro Iniesta, J. C., et al. 2007, Publications of the Astronomical Society of Japan, 59, 837
Paxman, R. G., Schulz, T. J., & Fienup, J. R. 1992, Journal of the Optical Society of America A, 9, 1072
Puschmann, K. G., Kneer, F., Seelemann, T., & Wittmann, A. D. 2006, Astronomy & Astrophysics, 451, 1151
Rees, D. E., Durrant, C. J., & Murphy, G. A. 1989, The Astrophysical Journal, 339, 1093
Richardson, W. H. 1972, Journal of the Optical Society of America, 62, 55
Ruiz Cobo, B. & del Toro Iniesta, J. C. 1992, The Astrophysical Journal, 398, 375
Scharmer, G. B. 2006, Astronomy & Astrophysics, 447, 1111
Scharmer, G. B., Henriques, V. M. J., Kiselman, D., & de la Cruz Rodríguez, J. 2011, Science, 333, 316
Skumanich, A. & Lites, B. 1991, in Solar Polarimetry, ed. L. J. November, 307–317
Socas-Navarro, H. 2003, Neural Networks, 16, 355
Socas-Navarro, H., Ruiz Cobo, B., & Trujillo Bueno, J. 1998, The Astrophysical Journal, 507, 470
Stenflo, J. 1968, Acta Univ.Lund. II, 2
Suematsu, Y., Tsuneta, S., Ichimoto, K., et al. 2008, Solar Physics, 249, 197
Title, A. & Rosenberg, W. A. 1981, Optical Engineering, 20, 815
Tsuneta, S., Ichimoto, K., Katsukawa, Y., et al. 2008, Solar Physics, 249, 167
van Noort, M., Rouppe van der Voort, L., & Löfdahl, M. G. 2005, Solar Physics, 228, 191
Vögler, A., Shelyag, S., Schüssler, M., et al. 2005, Astronomy & Astrophysics, 429, 335
von der Luehe, O. 1993, Astronomy & Astrophysics, 268, 374